

An “All Possible Steps” Approach to the Accelerated Use of Gillespie’s Algorithm

Azi Lipshtat

*Department of Pharmacology and Biological Chemistry,
Mount Sinai School of Medicine, New York, NY 10029**

Abstract

Many physical and biological processes are stochastic in nature. Computational models and simulations of such processes are a mathematical and computational challenge. The basic stochastic simulation algorithm was published by D. Gillespie about three decades ago [D.T. Gillespie, *J. Phys. Chem.* **81**, 2340, (1977)]. Since then, intensive work has been done to make the algorithm more efficient in terms of running time. All accelerated versions of the algorithm are aimed at minimizing the running time required to produce a stochastic trajectory in state space. In these simulations, a necessary condition for reliable statistics is averaging over a large number of simulations. In this study I present a new accelerating approach which does not alter the stochastic algorithm, but reduces the number of required runs. By analysis of collected data I demonstrate high precision levels with fewer simulations. Moreover, the suggested approach provides a good estimation of statistical error, which may serve as a tool for determining the number of required runs.

I. INTRODUCTION

Statistics and dynamics of stochastic systems have attracted a great deal of interest in many scientific fields, including physics, ecology, chemistry and biology.^{1,2,3}. The small and discrete number of reactant molecules in a cell leads to fluctuation-dominated dynamics. This type of dynamics appears to have important consequences in biology^{4,5,6}. Time evolution of such systems cannot be treated by standard continuous-time deterministic differential equations. For proper mathematical modeling of fluctuating systems, a probabilistic method is often necessary. Analytical solution of the probabilistic Master Equation is rarely available. In most cases the data is obtained from numerical simulations. There are several algorithms for simulating stochastic systems^{7,8,9,10,11,12,13,14,15}. However, all these methods are focused on the simulation algorithm, but not on efficient analysis of the output. A high precision level requires many runs of the algorithm. Here we propose a method of reducing the number of runs which are required for reliable estimation of moments. By doing so we reduce significantly the computational resources needed for stochastic simulations. We start with a general description of stochastic systems and the standard stochastic simulation algorithm. Our new analysis method is presented in Sec. II. In Sec. III we extend this approach to time course simulations. A discussion and summary are presented in Sec. IV.

A. Stochastic systems

A stochastic system consists of $N \geq 1$ molecular species, $\{S_1, \dots, S_N\}$, interacting through $M \geq 1$ chemical reactions $\{R_1, \dots, R_M\}$. The state of the system is defined by the molecular populations $\mathbf{X}(t) = (X_1(t), \dots, X_N(t))$. This is a random variable whose dynamics is determined by the reactions. For example, an occurrence of the reaction $S_1 + S_2 \rightarrow S_3$ changes the system's state from (x_1, x_2, x_3) to $(x_1 - 1, x_2 - 1, x_3 + 1)$. The stoichiometric vector of a reaction R_μ is denoted by ν_μ . The j th component of this vector is the discrete change in the population size of S_j resulting from a single occurrence of an R_μ reaction. Thus, the reaction takes the system from state \mathbf{x} to $\mathbf{x} + \nu_\mu$. Given that the state of the system at time t is $\mathbf{X}(t) = \mathbf{x}$, the probability that a reaction R_μ will occur during a time interval dt , is given by the *propensity function* $a_\mu(\mathbf{x})dt$. This function takes into account both the constant physical coefficients and the combinatorial factors which are state-dependent.

Given $\mathbf{X}(t) = \mathbf{x}$, the probability that the next reaction will be R_μ , and that it will occur in the time interval $[t + \tau, t + \tau + d\tau)$, is given by

$$p(\tau, \mu | t, \mathbf{x}) = a_\mu(\mathbf{x}) \exp(-a_0(\mathbf{x})\tau) \quad (1)$$

where

$$a_0(\mathbf{x}) = \sum_{\mu=1}^M a_\mu(\mathbf{x}). \quad (2)$$

The probability in Eq. 1 can be interpreted as a product of two independent probabilities: The probability to have R_μ as the next reaction is a_μ/a_0 , and the probability that next reaction will take place after about τ time units is $a_0(\mathbf{x}) \exp(-a_0(\mathbf{x})\tau)$.

B. Gillespie's algorithm

The standard algorithm for stochastic simulations was developed by Gillespie⁸. The algorithm follows a trajectory $(\mathbf{X}(t_0), \mathbf{X}(t_1), \dots, \mathbf{X}(t_{\text{final}}))$ of the system in the state space. The actual time interval between two successive reactions is a random variable distributed exponentially with average a_0^{-1} . Thus, the algorithm can be summarized as follows:

$$\text{Initialize } t = t_0, \mathbf{X} = \mathbf{X}(t_0). \quad (3)$$

While $t < t_{\text{final}}$

- calculate rates a_μ and a_0 for current state \mathbf{x} .
- draw τ from exponential distribution with mean a_0^{-1} .
- advance time from t to $t + \tau$.
- pick next reaction to be R_μ with probability a_μ/a_0 .
- update state \mathbf{x} to $\mathbf{x} + \nu_\mu$.

The algorithm is computationally expensive and intensive efforts have gone into making it more efficient^{9,10,11,12,13,14,15,16,17}. In his original work, Gillespie suggested two versions of the algorithm - the Direct method (DM) which is shown in 4 and the First Reaction method (FRM)⁸. Each of these was later optimized and found to be superior in certain cases^{9,10,17}. In addition to those exact algorithms, there are also some approximations which significantly accelerate the simulation performance, with minor loss of precision^{11,12,13,14,15,18}. The improvement is being done by either hybridization (simulate some of the species in a deterministic manner) or averaging over time intervals.

C. Data collection from simulations

Gillespie writes⁸:

“If it is desired to estimate any of the *moments* $X_i^{(k)}(t)$ of the grand probability function... then it will be necessary to make *several* simulation runs from time 0 to the chosen time t ... Any moment $X_i^{(k)}(t) \equiv \langle X_i^k \rangle_t$ may then be estimated directly as the average of the k th power of the numbers found for X_i at time t in these runs.”

The first moment is of interest in most cases, but higher moments may also be important. The second moment $\langle X_i^2 \rangle$ determines the variation, and combined moments $\langle X_i X_j \rangle$ give reaction rate of second order reactions. Even though estimation of the various moments is the typical aim of a stochastic simulation, one can hardly find explicit reference to the moment calculation in the literature. Very often the algorithm is described in great detail. The selection procedure for the next reaction and choice of data structure are elaborately explained. However, there is no mention of the translation of the resulting trajectory into moments. It is implicitly assumed that the statistics extraction process is straightforward. However, there may be different ways of performing the analysis. If a time course of a moment is needed, one has to define a set of time points. These are not the reaction times, which are different for any execution of the algorithm. The values of the moment at these time points are recorded and averaged over many runs of the algorithm. In the more common case, the focus is on steady state values of the moments, rather than on transients. It is assumed that there is a steady state distribution, and the probability $P(\mathbf{x})$ to find the system in state \mathbf{x} does not change in time. In that case, one should average not only over many runs but also over time within each run of the simulation. In practice, one has to calculate the sum

$$\sigma_j^k = \sum_i X_j^k(t_i) \tau_i, \quad (4)$$

where $X_j^k(t_i)$ is the value of the moment X_j^k at time point t_i in which the i th step (reaction) was taken, and $\tau_i = t_{i+1} - t_i$ is the time to the next reaction. After the simulation run is completed, the moment is given by

$$\langle X_j^k \rangle = \frac{1}{t_{\text{final}} - t_0} \sigma_j^k. \quad (5)$$

The moments can also be calculated using an indirect approach. In this approach, the stationary probability distribution $P(\mathbf{x})$ is measured during the simulation. Any moment can be deduced later from this distribution. Calculation of probabilities, rather than direct estimation of moments is advantageous in systems with multiple stable states, where the average moments obscure the existence of several possible stable states^{19,20}. The probability distribution $P(\mathbf{x})$ is calculated as the fraction of time that the system spent at state \mathbf{x} . If the system remained in state \mathbf{x} for time τ , then the non-normalized probability $q(\mathbf{x})$ has to be updated:

$$q(\mathbf{x}) = q(\mathbf{x}) + \tau, \quad (6)$$

where $q(\mathbf{x})$ was initialized to zero at the beginning of the simulation. The probability distribution $P(\mathbf{x})$ is given at the end of the simulation by

$$P(\mathbf{x}) = \frac{q(\mathbf{x})}{t_{\text{final}} - t_0}. \quad (7)$$

We refer to this method of calculation as the “Trajectory Following” approach.

To make sure that states are sampled according to the correct distribution, one has to start collecting the data not immediately at the beginning of the simulation, but after some equilibration time. During that time the system should arrive its stationary probability distribution. Including transient in the collected data may change the calculated moments, and require much longer simulations, so that this change will be negligible.

Either in the direct or indirect methods, one has to repeat the simulation many times and average the results over a large ensemble of simulations. The reason is that in the typical case there are states whose probability is low, which are visited during a standard run of a simulation for short durations and a very small number of times, if at all. Including the effect of these rare events on the average moment requires a large enough number of visits, so that the portion of time in which the system is in these states would be proportional to their probabilities. The error in stochastic simulations decreases as $1/\sqrt{N}$, where N is the number of runs^{21,22}. Thus high precision requires large N and longer running time.

II. THE “ALL POSSIBLE STEPS” APPROACH

Here we suggest a novel approach for calculating the probabilities $P(\mathbf{x})$, namely the “*All Possible Steps*” (APS) method. This method requires a smaller number of runs for a given

precision, and thus makes the whole simulation more efficient. The “*All Possible Steps*” method is based on calculating the probability of visiting states without actually visiting them. At any step in the simulation, the probability for occurrence of reaction R_μ and moving the system from \mathbf{x} to $\mathbf{x} + \nu_\mu$ is given by $a_\mu(\mathbf{x})/a_0$. In the standard methods, we choose one of the possible reactions and update the moment or the probabilities accordingly. In the APS approach, we update the collected data by considering *all* possible steps, each with a weight proportional to its probability. At the i th step, the system is at state $\mathbf{x} = \mathbf{x}(t_i)$. There are several possibilities for the next state. There is a probability $a_\mu(\mathbf{x})/a_0(\mathbf{x})$ to choose reaction R_μ and change the state of the system to $\mathbf{x} + \nu_\mu$. If this happened, the system will be at that state for a duration τ_μ which is taken from an exponential distribution with average $(a_0(\mathbf{x} + \nu_\mu))^{-1}$. Thus, we consider this step as if it was done and update the (non normalized) probability by adding $(a_\mu(\mathbf{x})/a_0(\mathbf{x})) \tau_\mu$. For the purpose of statistics collection, we consider virtually *all* possible reactions and update *all* probabilities respectively. Then we choose one of them and update the state accordingly. To keep the ratio between probabilities correct, we use the same random number when calculating the duration τ_μ for each of the possible steps. When the run is over, the probabilities have to be normalized. Since we considered many steps simultaneously, the normalization is not done by the total running time. One has to divide the probabilities by their sum, so that the sum of all probabilities would be equal to unity. The differences between data collection in the standard Trajectory Following approach and in the APS method for steady state simulations are summarized in Table I.

A. Example: Protein dimerization

As an example we take a schematic process of protein dimerization, which includes three reactions: production, degradation and dimerization.



In this example, the Master equation is solvable analytically. In a different context, it was shown^{23,24} that the probability to have n copies of S_1 is

$$P_{\text{exact}}(n) = \frac{2^{\frac{1}{2}(\frac{k_2}{k_3}-1)}}{n!} \left(\sqrt{\frac{k_1}{k_3}} \right)^n \frac{I_{k_2/k_3+n-1} \left(2\sqrt{k_1/k_3} \right)}{I_{k_2/k_3-1} \left(2\sqrt{2k_1/k_3} \right)}, \quad (9)$$

and the average copy number is

$$\langle S_1 \rangle = \sqrt{\frac{k_1}{2k_3}} \frac{I_{k_2/k_3} \left(2\sqrt{2k_1/k_3} \right)}{I_{k_2/k_3-1} \left(2\sqrt{2k_1/k_3} \right)}. \quad (10)$$

where $I_x(y)$ is the modified Bessel function. A state in this system is defined by the number n of S_1 copies. In the Trajectory Following approach, at any state n , only the probability $p(n)$ of the current state will be updated

$$q(n) \longrightarrow q(n) + \tau. \quad (11)$$

However, in the APS method we consider all possible steps, even those who were not chosen for actual move. Hence, after the system moved to state n , we update all possible states for the next step, namely $n - 2, n - 1$ and $n + 1$, with appropriate weights:

$$\begin{aligned} q(n-1) &\longrightarrow q(n-1) + \tau_{n-1} a_{n \rightarrow n-1} / a_0 \\ q(n+1) &\longrightarrow q(n+1) + \tau_{n+1} a_{n \rightarrow n+1} / a_0 \\ q(n-2) &\longrightarrow q(n-2) + \tau_{n-2} a_{n \rightarrow n-2} / a_0, \end{aligned} \quad (12)$$

where τ_n is the random time duration the system would spend in state n , had this been the next step, and $q(n)$ is the probability to be at state n , before normalization.

Updating several probabilities at any step requires of course more computations. However, since these are only standard arithmetic operations, there is no significant overhead in terms of running time. The most expensive steps in a stochastic simulation are the generation of pseudo random numbers and taking decisions based on those numbers. In the APS approach there is no need to have any extra random numbers and the update of probabilities is simple, so the extra running time is negligible. Furthermore, the whole analysis can be completely decoupled from the trajectory production process. One can produce the trajectory using the standard algorithm, and at a later time make the analysis and calculate the moments using the APS approach.

Accuracy of simulation can be measured by comparison with the exact solution (9). The error in calculation of the first moment $\langle S_1 \rangle$ is shown in Fig. 1. The relative deviation

$$\varepsilon = \frac{|\langle S_1 \rangle_{\text{exact}} - \langle S_1 \rangle_{\text{sim}}|}{\langle S_1 \rangle_{\text{exact}}} \quad (13)$$

was calculated using both the Trajectory Following approach and the APS method. In both cases the error decreases as $1/\sqrt{N}$ where N is the number of runs, as expected. However, the error in the APS approach is in average about 4 times smaller than the error in the standard Trajectory Following method. Thus, to obtain a specific precision, one needs about 16 fold more simulations in the standard method than in the APS method. For example, for $\varepsilon = 10^{-3}$ (precision of 0.1%), average over 874 simulations is required at the standard method, whereas 47 runs are enough in the APS approach (values obtained from interpolation of $N(\varepsilon)$. See thin lines in Fig. 1). This is almost a 19 fold improvement.

The distribution of errors in the APS method came out to be much narrower than that of the Trajectory Following approach. To show this, I ran 1000 simulation sessions. Each session continued until a precision level of $\varepsilon = 10^{-3}$ was obtained. In the APS approach, 76% of the simulations ended with an estimation for $\langle S_1 \rangle_{\text{sim}}$ which is within $\pm 1\%$ of the exact value $\langle S_1 \rangle_{\text{exact}}$ (and about 9% within $\pm 0.1\%$), whereas in the Trajectory Following method only 22% were within the 1% vicinity of $\langle S_1 \rangle_{\text{exact}}$ (and 2% in the 0.1% range).

III. TIME COURSE SIMULATIONS

In principle, it is easy to use the same approach for time course simulations. In this case, one should update only the probabilities at those time points which fall into the interval $[t_i, t_i + \tau_\mu)$. The probability to find the system at these time points in state $\mathbf{x} + \nu_\mu$ is increased by $a_\mu(\mathbf{x})/a_0$. The normalization by the end of the simulation would be such that the sum of probabilities at each time point will be equal to one. This procedure is summarized in Table II.

In practice, it turns out that for this type of simulations, the improvement of the APS method is lower than in the steady state case. In time course simulations, the rare states are important because they lead to different trajectories. In the APS method we do not compute these hypothetical trajectories, and thus the resulting dynamics is not significantly different from that of the standard algorithm.

Nevertheless, there are cases in which the APS method is useful for time course simulations. In many systems, there is a separation between “slow” and “fast” reactions or species. In fact, some stochastic simulation algorithms are based on this separation^{14,25,26}. Highly reactive intermediate species, such as enzyme-substrate complexes, usually obey the quasi-steady-state assumption (known also as the pseudo-steady-state assumption). This means that the typical time scale for a significant change in populations of these species is much longer than the sampling time. As a result, at any given time, these species can be considered to be close enough to steady state. Thus, using the APS method for simulating the stochastic system does not affect the results of the “slowly reacting” species, but definitely improves the statistics as far as “fast” species are concerned. We demonstrate this in a slow dimerization reaction of fast species (example 6.2 of Ref. 26). The system consists of three species, out of which two are reactive:



where the reaction rate coefficients are given by $(k_1, k_{-1}, k_2) = (200, 1, 1)$ and at initial time $[S_1] = 0, [S_2] = 100, [S_3] = 0$. The net change in the S_1 population size is close to zero, and thus S_1 is in quasi-steady-state. In Fig. 2 we present the dynamics of $[S_1]$, as obtained from standard averaging and from using the APS method. Even though this is not a real steady state, the change in the population size is slow.

A. The virtual population approach

The concept of considering virtual steps can be beneficial in the context of time course simulations in an additional way. In most systems there are products which are not involved in any reaction as reactants. S_3 is such a species in our last example. Since these species are not reactive, their population size has no effect on the system dynamics. Thus, there is no need to follow the real discrete population size. Instead, one can update the *virtual* population size at any step of simulation, whether or not a real production reaction took place in that step. The update should be by the reaction rate (as obtained from the discrete system state) multiplied by the time step size. In our example, we increased the population

size of S_3 at any step by $k_2[S_1]([S_1] - 1)\tau$, where $[S_1]$ is the discrete population size of S_1 and τ is the time step, calculated by the algorithm as function of the system's state. Mathematically, this approach is equivalent to calculating $[S_3]$ by

$$[S_3(t)] = [S_3(0)] + \int_0^t \frac{d[S_3]}{dt'} dt'. \quad (15)$$

In Fig. 3 (a) we present the population size of S_3 in a single run as calculated by the Trajectory Following approach and by the aforementioned method. Since the population size is updated at any step, and not only in the rare events of real dimer production, the curve is much smoother and not fluctuating. The variation between different runs in this method is smaller as well, as shown in Fig. 3 (b).

IV. DISCUSSION AND SUMMARY

We have proposed a new method of analysis for stochastic simulations data. The “*All Possible Steps*” method gives a precise estimation of moments of the variables, which is based on all possible steps of the simulation. Since the information about those steps and their probabilities is calculated in any case throughout the simulation, the *All Possible Steps* method does not require many additional computation resources. On the other hand, reduction in fluctuation size and variation enables one to achieve reliable statistics with fewer runs of the algorithm. The method is applicable mainly for simulations under steady state condition, but it can be extended also to some transient cases.

Furthermore, the availability of two methods for moment calculation can be used to determine the number of runs required for any desired level of precision. In most cases, there is no way of estimating how many runs are needed to obtain a given precision. As a result, in the typical case a large number of simulations are being performed, many more than really required. The existence of two estimation methods for the moments provides a tool for the estimation of statistical error. The relative difference between the methods is given by

$$\eta = \frac{|M_{\text{TF}} - M_{\text{APS}}|}{M_{\text{APS}}}, \quad (16)$$

where M_{TF} and M_{APS} are estimations for the moment as calculated by the Trajectory Following and APS methods, respectively. η is a good estimator for the average error. Instead of defining in advance the number of simulation runs, one can run the simulation

until η is smaller than a pre-defined threshold. This way, we avoid unnecessary simulations. In Fig. 4, we show that η decreases like $1/\sqrt{N}$, where N is number of runs, as expected for the statistical error.

There are acceleration methods which are based on approximation to the exact Trajectory Following algorithm. This is not the case in the APS approach, which is mathematically equivalent to averaging over many standard trajectories. In the Trajectory Following algorithm, the probability of a certain state is calculated directly as the fraction of time that the system spent in that state. In the *All Possible Steps* approach we manipulate the data obtained from the actual trajectory to reflect virtual averaging over some hypothetical trajectories as well. However, we consider only the first step of each of those trajectories. Since the actual trajectory is determined by the original algorithm, the number of times a state \mathbf{x} is being visited is not affected by considering the hypothetical steps. As a results, the number of occurrences of a reaction $\mathbf{x} \rightarrow \mathbf{x} + \nu$ will be calculated properly, and so is the time of staying in the new state $\mathbf{x} + \nu$.

It may be possible to extend the method to include more than one hypothetical reaction. In that case one should consider not only all possible reactions $\mathbf{x} \rightarrow \mathbf{x} + \nu$, but also all possible paths $\mathbf{x} \rightarrow \mathbf{x} + \nu_1 \rightarrow \mathbf{x} + \nu_2$. However, this requires more calculations, especially if the number of reactions is large.

The approach described here does not produce trajectories which are different from those resulting from the standard algorithms. It does not affect the stochastic simulation algorithm itself, but only the way data are collected. Thus, it can apply to many variants of the simulation algorithm. The improvement achieved by using this method comes *in addition to*, and not instead of, other algorithm-based accelerations. Furthermore, incorporating the APS approach into some of the accelerating algorithms may enhance the effect of using the APS method, due to considering possible trajectories rather than possible single steps. We leave this for future research. It is expected that more sophisticated versions will be developed. Progress in the two complementary approaches - reducing the number of required runs on the one hand, and increased efficiency of every single run on the other hand - will make simulations of larger systems computationally inexpensive.

Acknowledgments

I am grateful to Ravi Iyengar (MSSM) for his help. I wish to thank also Charles Peskin (NYU), Fernand Hayot, Eric Sobie and Padmini Rangamani (MSSM) for careful reading of the manuscript and for important comments.

* E-mail: Azi.Lipshtat@mssm.edu

- ¹ N.G. van Kampen, *Stochastic Processes in Physics and Chemistry* (North-Holland, Amsterdam, 1981).
- ² M.S. Bartlett, *Stochastic Population Models in Ecology and Epidemiology* (John Wiley & Sons, Inc., NY, NY, 1961).
- ³ S. Karlin and H.M. Taylor, *An Introduction to Stochastic Modeling, Third Edition* (Academic Press, San Diego, 1998).
- ⁴ H. H. McAdams and A. Arkin, *Proc. Natl. Acad. Sci. US* **94**, 814 (1997).
- ⁵ H. H. McAdams and A. Arkin, *Trends Genet.* **15**, 65 (1999).
- ⁶ M.B. Elowitz, A.J. Levine, E.D. Siggia, and P.S. Swain, *Science* **297**, 1183 (2002).
- ⁷ D.T. Gillespie, *J. Comput. Phys.* **22**, 403 (1976).
- ⁸ D.T. Gillespie, *J. Phys. Chem.* **81**, 2340 (1977).
- ⁹ M.A. Gibson and J. Bruck, *J. Phys. Chem.* **104**, 1876 (2000).
- ¹⁰ Y. Cao, H. Li and L. Petzold, *Comp. Biol. and Chem.* **30**, 39 (2006).
- ¹¹ D.T. Gillespie, *J. Chem. Phys.* **115**, 1716 (2001).
- ¹² E.L. Haseltine and J.B. Rawlings, *The Journal of Chemical Physics* **117**, 6959 (2002).
- ¹³ D.T. Gillespie and L.R. Petzold, *J. Chem. Phys.* **119**, 8229 (2003).
- ¹⁴ C.V. Rao and A.P. Arkin, *The Journal of Chemical Physics* **118**, 4999 (2003).
- ¹⁵ Y. Cao M. Rathinam, L.R. Petzold and D.T. Gillespie, *The Journal of Chemical Physics* **119**, 12784 (2003).
- ¹⁶ L. Lok, *Nature Biotechnology* **22**, 964 (2004).
- ¹⁷ J.M. McCollum, G.D. Peterson, C.D. Cox, M.L. Simpson and N.F. Samatova, *J. Comput. Phys.* **206**, 395 (2005).
- ¹⁸ H. Resat, H.S. Wiley and D.A. Dixon, *J. Phys. Chem. B* **105**, 11026 (2001).

¹⁹ A. Lipshtat, A. Loinger, N.Q. Balaban and O. Biham, *Physical Review Letters* **96**, 188101 (2006).

²⁰ A. Loinger, A. Lipshtat, N.Q. Balaban and O. Biham, *Physical Review E* **75**, 021904 (2007).

²¹ K. Binder and D.W. Heermann, *Monte Carlo Simulation in Statistical Physics*, (Springer, NY, NY, 1997).

²² G.S. Fishman, *Monte Carlo* , (Springer, NY, NY, 1996).

²³ O. Biham and A. Lipshtat, *Phys. Rev. E* **66**, 056103 (2002).

²⁴ N.J.B. Green, T. Toniazzo, M.J. Pilling, D.P. Ruffle, N. Bell and T.W. Hartquist, *Astron. Astrophys.* **375**, 1111 (2001).

²⁵ J. Puchalka and A.M. Kierzek, *Biophys. J.* **86**, 1357 (2004).

²⁶ Y. Cao, D. Gillespie and L. Petzold, *J. Comput. Phys.* **206**, 395 (2005).

	Trajectory Following	All Possible Steps
Initialization before the simulation		$Q(\mathbf{x}) = 0$
Initialization before every run		$q(\mathbf{x}) = 0$
Update (every iteration)		<p>calculate τ</p> <p>$t = t + \tau$</p>
	$q(\mathbf{x}) = q(\mathbf{x}) + \tau$	<p>For each reaction R_μ</p> <p>$q(\mathbf{x} + \nu_\mu) = q(\mathbf{x} + \nu_\mu) + \frac{a_\mu}{a_0} \tau_\mu$</p>
		pick reaction R_μ and update state to $\mathbf{x} + \nu_\mu$
Normalization after every run	$Q(\mathbf{x}) = Q(\mathbf{x}) + \frac{q(\mathbf{x})}{t_{\text{final}} - t_0}$	$Q(\mathbf{x}) = Q(\mathbf{x}) + \frac{q(\mathbf{x})}{\sum_{\mathbf{x}} q(\mathbf{x})}$
Normalization after N runs		$P(\mathbf{x}) = Q(\mathbf{x})/N$

TABLE I: Comparison between the standard method for probability calculation and the *All Possible Steps* method for steady state simulations. $q(\mathbf{x})$ and $Q(\mathbf{x})$ are probability distributions before normalization.

	Trajectory Following	All Possible Steps
Initialization before first run	for each time point t_i , $q(\mathbf{x}, t_i) = 0$	
Update (every iteration)	calculate τ	
	For each t_i s.t. $t < t_i \leq t + \tau$ $q(\mathbf{x}, t_i) = q(\mathbf{x}, t_i) + 1$ $t = t + \tau$	$t = t + \tau$ For each reaction R_μ For each t_i such that $t < t_i \leq t + \tau_\mu$ $q(\mathbf{x}, t_i) = q(\mathbf{x}, t_i) + a_\mu/a_0$
	Pick reaction R_μ and update state to $\mathbf{x} + \nu_\mu$	
Normalization after N runs	$P(\mathbf{x}, t_i) = \frac{q(\mathbf{x}, t_i)}{N}$	$P(\mathbf{x}, t_i) = \frac{q(\mathbf{x}, t_i)}{\sum_{\mathbf{x}} q(\mathbf{x}, t_i)}$

TABLE II: Comparison between the standard method for probability calculation and the APS method for time course simulations. $q(\mathbf{x})$ is the probability distribution before normalization.

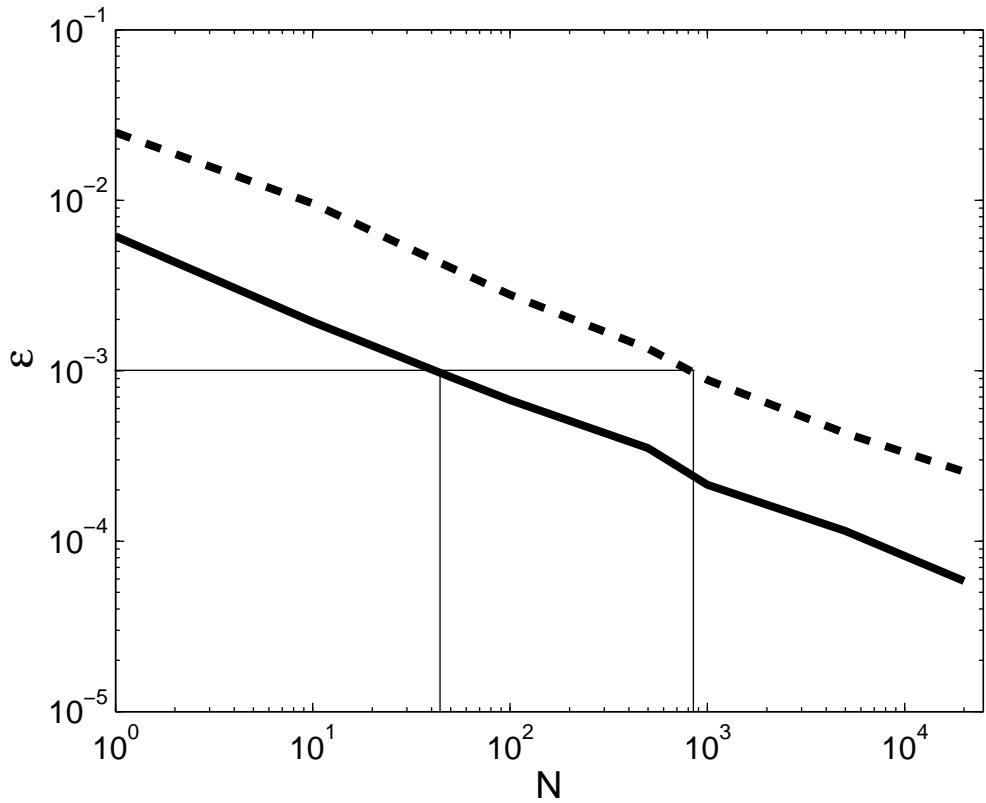


FIG. 1: Error measurement (Eq. 13) as function of N , number of runs. ε was calculated for the *same* set of simulations using both the Trajectory Following approach (dashed line) and the APS method (solid line). The rates are $k_1 = 5, k_2 = 2, k_3 = 2$. Initial condition is $S_1(0) = S_2(0) = 0$. Each simulation ran for 10 time units, which is enough for attaining steady state, and then statistics was collected for 150 time units ($t_0 = 10, t_{\text{final}} = 160$). For each value of N the error was averaged over an ensemble of 100 realizations of N simulations. The thin lines demonstrate the 19 fold improvement in case of $\varepsilon = 10^{-3}$. See text for details.

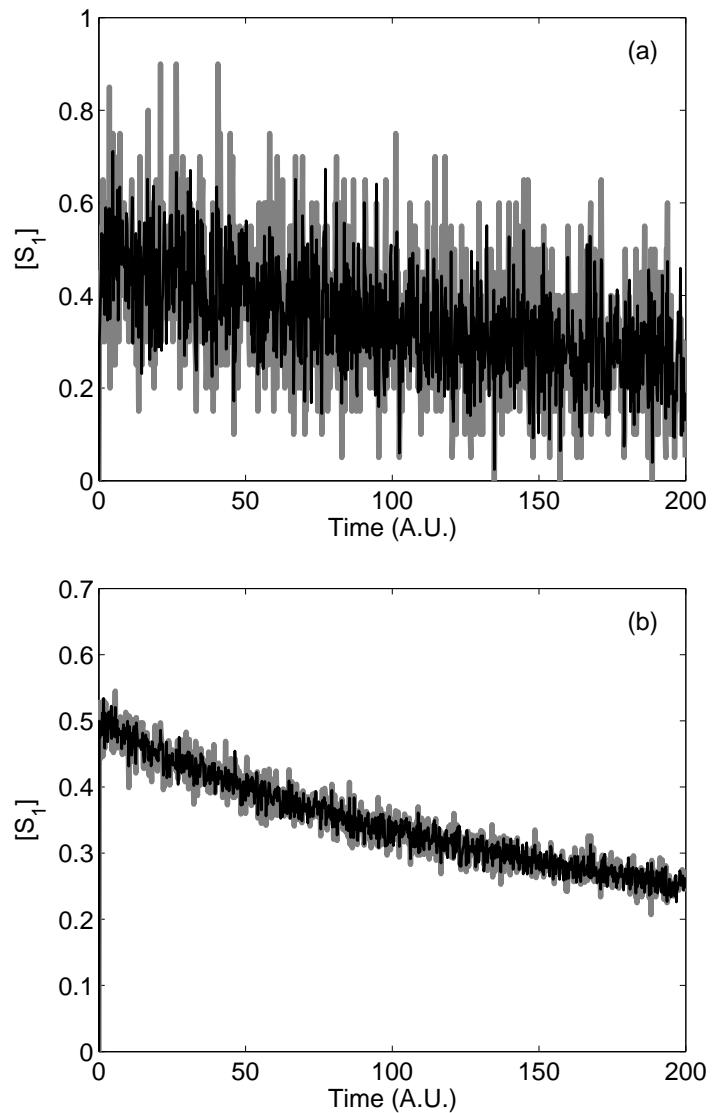


FIG. 2: Dynamics of $[S_1]$ for reactions 14. Both Trajectory Following approach (gray) and APS approach (black) show similar dynamics, but the Trajectory Following results show more fluctuations. Averaged over (a) 20 simulations, (b) 1000 simulations.

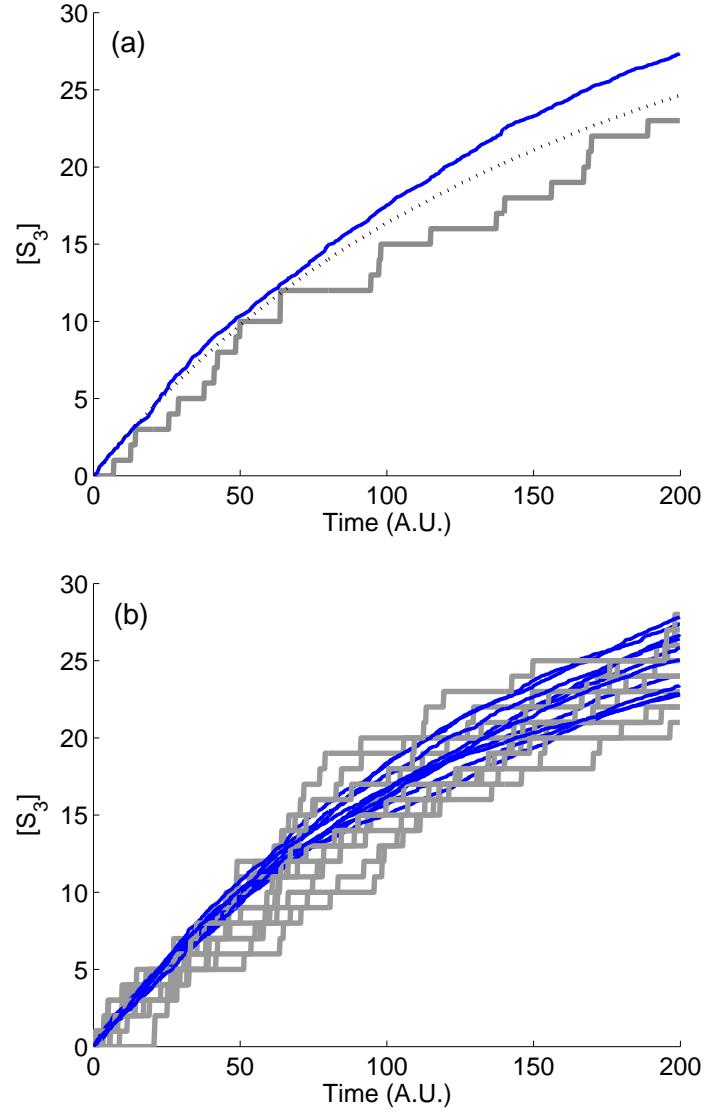


FIG. 3: Dynamics of $[S_3]$ in (a) one single run and (b) 10 single runs, using Trajectory Following approach (gray step-like lines) and virtual population approach (dark smooth curves). The virtual population approach yields smoother curves with smaller variation. The dotted line in (a) is an average over 200 runs, where both the methods are indistinguishable.

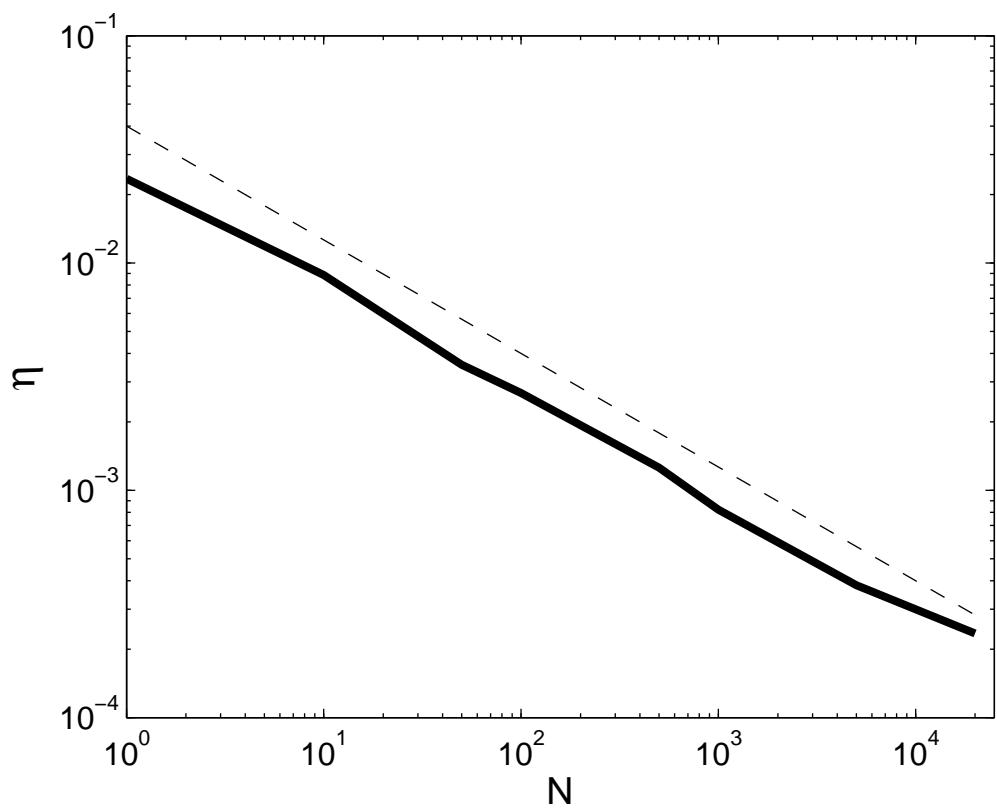


FIG. 4: Estimation for the error as function of number of simulations. The dashed line is proportional to $1/\sqrt{N}$. All simulation parameters are as in Fig. 1.